arXiv:1611.03780v1 [cs.SI] 11 Nov 2016

GeoCUTS: Geographic Clustering Using Travel Statistics

David Rolnick* Kevin Aydin† Shahab Kamali‡ Vahab Mirrokni§ Amir Najmi¶

November 14, 2016

Abstract

Web-based services often run experiments to improve their products. To carry out an effective experiment and evaluate the results appropriately, there must be a control group and at least one treatment group. Ideally all of these groups are disjoint, so that each user is given a specific treatment. Using geographical locations as units of experimentation is desirable because this does not require tracking individual users or browser cookies. With the popularity of mobile devices, a user may issue queries from multiple geographical locations. Hence, to be used as units of experimentation, geographical partitions should be chosen in a way that reduces transit between regions.

The strategy of clustering users by region is common in advertising. Designated marketing areas (DMAs) are specifically designed for this purpose. However, DMAs are restricted to the US and their granularity is inflexible (there are around two hundred in total). Moreover, they are built based on population density - one DMA per metropolitan area - rather than mobile movement patterns.

In this paper, we present GeoCUTS, an algorithm that forms geographical clusters to minimize movement between clusters while preserving rough balance in cluster size. We use a random sample of anonymized mobile user traffic to form a graph representing user movements, then construct a geographically coherent clustering of the graph. We propose a statistical framework to measure the effectiveness of clusterings and perform empirical evaluations showing that the performance of GeoCUTS is comparable to hand-crafted DMAs with respect to both novel and existing metrics. GeoCUTS offers a general and flexible framework for conducting geo-based experiments in any part of the world.

1 Introduction

Large-scale online services routinely conduct live experimentation to improve their products. As described in [17], browser cookies are the standard unit of analysis for experiments run by many web services. Typically, cookies are randomly selected into disjoint treatment and control groups and then subjected to different treatments (e.g. a different color background). The statistical and practical significance of metric differences between the two groups of cookies are important factors in the decision whether or not to launch the experimental treatment to all users.

While this cookie-based approach has been used extensively in the industry, it has some important limitations when measuring long-term effects on users. As detailed in [10], the fundamental problem is that

> "[cookies]...are a poor proxy for users. A cookie is simply an anonymous id attached to a browser and a device. Users can clear their cookies whenever they want, and they frequently use multiple devices and multiple browsers."

The authors further observe that

*Massachusetts Institute of Technology, Cambridge, MA USA, drolnick@math.mit.edu
†Google, Inc., New York, NY USA, kaydin@google.com
‡Google, Inc., Mountain View, CA USA, kamali@google.com
§Google, Inc., New York, NY USA, mirrokni@google.com
¶Google, Inc., Mountain View, CA USA, amir@google.com

> "Using a signed-in id may seem to mitigate this issue, but users can have multiple sign-ins and many queries are made while signed-out."

Thus the effect measured by a cookie experiment is diluted by the fact that a user may be in the treatment group on one device or browser and in the control group in another.

A different approach to treating *whole users* in an experiment is to experiment on randomized geographical regions. This is a form of *cluster sampling*. Assuming a user remains within one such region, the user will receive consistent experimental treatment regardless of which device or browser she uses. The use of geo regions is common in advertising, with Nielsen's DMA (Designated Market Area)® regions being the best known. These have been used to run randomized *geo experiments* for measuring the effectiveness of internet advertising [21].

DMAs have the advantage of being well-established as a means of subdividing a user population. However, they were created by Nielsen to correspond to television audiences in the US. Their partitioning is not necessarily optimal for internet usage. Furthermore, DMAs are restricted to the US without a direct international equivalent, and their granularity is fixed, with approximately 200 across the US.

We should also mention that the US Office of Management and Budget defines 381 Metropolitan Statistical Areas (MSAs) for the US. Each MSA is an economically integrated set of adjacent counties with at least one urban core area of 50,000. Since MSAs suffer from the limitations of DMAs noted above, we do not discuss them further in this paper.

The increase in mobile device usage poses an additional problem for the use of DMAs (and MSAs) as the unit of experimentation. This is because users can move into and out of a region. By doing so, they cause *leakage* between the treatment and control arms of the experiment. The growth of mobile devices has exacerbated the problem, though the same user can also use a single non-mobile device (e.g. laptop) in different geo regions or use two different fixed devices (e.g. one at work and one at home).

Ideally, geographic partitions should be based not on television viewership or population density but on the distribution and movement of users of the product in question. It would be useful to define these geo regions to strike a balance between the following design criteria:

1. Regions should be large enough to minimize leakage during the experiment (the probability of a user issuing queries from more than one region).

2. Regions should be numerous.

3. Regions should be relatively balanced in size.

#1 above helps us ensure users in a geo region will receive consistent treatment during experimentation, while #2 and #3 allow us to make assumptions of exchangeability necessary to draw inferences from an experiment. Even though we do not require perfection, each of these requirements is necessary if geo regions are to be a better basis for experimentation than cookies.

While treating whole users is the primary motivation to develop a geo-based experimental framework, there is an important secondary motivation as well. When experiments are run with large, noticeable changes (e.g. a new feature or significant redesign of an existing feature) it is often the case that users influence each other by word of mouth. In such a situation, part of the treatment administered to a randomly assigned group of users may leak out to affect the control group according to the patterns of social networks. To fully account for these effects would require social network information and is beyond the scope of this paper; however, such networks tend to be biased towards geographically local connections [4]. Thus we expect that our geo partitions, while designed to minimize splitting single users, are also somewhat well-suited to mitigate secondary effects from interaction between treatment and control groups.

In this paper, we present an algorithm, Geographic Clustering Using Travel Statistics (GeoCUTS) designed to optimize the three above criteria. We perform a comprehensive evaluation to study the impact of different design choices on GeoCUTS and also compare its performance against alternatives. We consider various massive-scale datasets to evaluate our algorithm under different circumstances, and consider both novel and existing metrics to compare GeoCUTS with alternative algorithms with respect to the balance

and leakage between clusters. The results of our experiments suggest that the performance of GeoCUTS is equal to or surpasses that of hand-designed regions. Moreover, our automatic approach is much simpler and easier, provides extra flexibility in controlling the number of clusters, and extends to countries besides the US.

The contributions of this paper are as follows:

- Modeling user movement patterns as a graph using geographic information.
- Designing the GeoCUTS algorithm to partition this graph into regions.
- Describing a statistical framework and introducing metrics to evaluate the quality of the clusters.
- Showing through a comprehensive evaluation that GeoCUTS matches or outperforms the output of expert annotation while allowing flexibility essential to randomized geo experiments.

2 Prior results

Many authors have considered the problem of causal inference on networks, where treatment is imposed at the level of nodes and interference is assumed to occur between nodes that are sufficiently close. For instance, in a social network, an individual's response may be influenced both by their own treatment and by that of their friends or friends of friends. Ugander et al. [20] and Gui et al. [9] design low-variance estimators for these cases, relying upon clusters within the network. Eckles et al. [7] demonstrate the effectiveness of *network bucket testing*, randomly assigning treatments to different clusters, which gives a natural low-variance estimator. This approach is analyzed further by Backstrom and Kleinberg [3] and by Katzir et al. [11], who consider making clusters using weighted random walks on the network.

In this paper, we build upon existing literature for finding clusters efficiently within a large network. Here we focus on a variant known as *balanced partitioning*. In balanced partitioning, our goal is to find a set of clusters of almost equal size and to minimize the total weight of edges that cross clusters (i.e., minimize the cut). This is a challenging problem that is computationally hard even for medium-sized graphs [1] as it captures the problem of graph bisection [8].

Balanced graph partitioning is not only an NP-hard problem, no constant approximation algorithm is known for it and solving this problem even for graphs with hundreds of nodes is not easy. In fact, the logarithmic approximation algorithms for this problem are based on solving linear programming and semi-definite programming relaxations for these problems. These LP and SDP relaxations are hard to solve for graphs with thousands of nodes, and we aim to solve this problem for much larger graphs. As a result, we need to settle for heuristic algorithms that can be implemented in a distributed manner.

While the topic of large-scale balanced graph partitioning has attracted significant attention in the literature [5, 6, 18, 19, 15], many prior authors have studied large-scale but non-distributed solutions to this problem. The need for distributed algorithms has been observed by several practical and theoretical research papers [2, 18, 19]. Here we summarize some related work in this area.

Zhu and Ghahramani [22] introduced the approach of label propagation. According to this method, clusters are initialized with known labels, after which these labels are propagated outward through the other nodes according to the transition probabilities of a Markov chain. The approach of label propagation was generalized to balanced clusters in the work of Ugander and Backstrom [19]. These authors formulated the swapping of nodes between clusters as a linear program which only depends on the number of clusters therefore allowing its implementation on very large graphs. More recently, Aydin et al. [2] achieve a scalable clustering in large networks by embedding the nodes along a line and use this embedding in future optimization steps. This approach has been proved to be effective for highly connected networks and expanders such as social networks [2].

Our work in this paper concentrates upon clustering with a geographically defined network, in contrast to the focus upon social networks. Recent work suggests that for graphs that have a geographic nature, applying the idea of "natural cuts" [5, 6] is effective in solving the balanced partitioning problem. Inspired by this heuristic, we develop a distributed implementation, which we detail further in Section 3.2.

3 Algorithm

We now present the GeoCUTS algorithm. The input is a set of locations from which a single user has issued queries. The algorithm proceeds in two phases. In Phase 1, we build a graph from the given data by discretizing locations and assigning edges between nodes for which there is frequent transit (see Algorithm 1). This algorithm is designed to be run massively in parallel. Each step in the computation can be distributed among individual vertices simply with knowledge of their neighborhood. In Phase 2, we find a clustering of this graph by applying a geographic clustering algorithm that combines recently developed techniques.

3.1 Phase 1: Graph building

```
for user in users do
    for A in locations(user) do
        a = num_visits(user, A);
        increase_weight(A, √a);
    end
    for A, B in locations(user) do
        a = num_visits(user, A);
        b = num_visits(user, B);
        if is_local(AB) then
            increase_weight(AB, √ab);
        end
    end
end
for A in nodes do
    normalize_weight(A);
end
for A,B in nodes do
    normalize_weight(AB);
end
```

Algorithm 1: Phase 1 of the GeoCUTS algorithm.

Discrete locations. The first step in building our graph is location discretization. While given user locations are assumed to be on a continuum, in real-world settings accuracy is finite. Therefore, we round each location to the nearest gridpoint in a lattice, where the width of the lattice may be specified; a coarser lattice yields a faster but less precise algorithm. We have chosen to discretize to a grid, rather than for instance to the nearest city, as described in Ugander and Backstrom [19]. The motivation is that a grid has a natural geometrical structure that we will utilize in our algorithm, and in addition our method is applicable to location data from users in rural areas or along transit corridors.

Node weights. We define a graph for which the nodes are the gridpoints above; in the following discussion we shall identify each node with the corresponding range of user locations. The weight of a node is a measure of the number of user visits to that location; specifically, if a user visits node A a total of a times, then the user's influence upon A is $\sqrt{a}$:

$$\text{weight}(A) = \sum_{\text{user } u} \sqrt{\#\text{ visits } u \text{ to } A}.$$

We use $\sqrt{a}$ as a slight normalization. Alternatively, it would be possible to increment the node's influence by a itself; however, some individual users may have much more location data available than others, and this would bias the graph strongly towards these users to the exclusion of others. At the opposite extreme, one might assign a fixed influence to each user and distribute this proportionally among the weights of the nodes

visited; however, this would bias the graph towards users for whom little location data is known, increasing the susceptibility to noise.

Edge weights. The edges in this graph correspond to the intensity of transit between nodes. Specifically, if a user visits node A a total of a times, and node B a total of b times, then the user's influence upon edge AB shall be $\sqrt{ab}$:

$$\text{weight}(AB) = \sum_{\text{user } u} \sqrt{(\#\text{ visits } u \text{ to } A) \cdot (\#\text{ visits } u \text{ to } B)}.$$

We use the geometric mean to increment edge weights since it is minimized when either endpoint is visited seldom, and is maximized when the endpoints are visited equally.

Sparsity. We retain only certain edges within the graph, specifically those for which the geographical distance between the two nodes is less than a given parameter. The reason for this trimming of edges is twofold: Firstly, it greatly reduces the size of the graph, allowing the number of edges to be linear, instead of at worst quadratic, in the number of nodes. Secondly, edges between nodes at great geographical distance can prove actively counterproductive to the algorithm. This is because the majority of user transit is geographically local, and therefore the desired output of the algorithm consists of clusters that are largely geographically connected. While there may be a relatively large amount of transit between, say, San Francisco and New York City, it would be a suboptimal to assign these cities to the same cluster. Therefore, any edge weight assigned to such long-distance connections would merely add noise to the objective function of our algorithm and make it more challenging to find the optimal clustering.

Normalization. As a final and critical step, we renormalize the weights in the graph. We consider the cases of renormalizing the weight x to $\sqrt{x}$ and to $\log(x)$. As shall be discussed further in Section 5, we observed best results after log-normalizing both node and edge weights, though the normalization of nodes was most critical. Our motivation is that we are more concerned with finding a natural clustering than with the exact balance of the pieces. Our clustering algorithm is designed to balance within hard limits, which would generally yield clusters constrained by size rather than graphical structure. For example, New York City, which contains many users, might be divided into many small clusters, irrespective of user traffic between them. By normalizing node weights, we ensure a condition of "soft balance" in which natural clusters are allowed to vary in size with a preference for balance.

3.2 Phase 2: Graph clustering

As a crucial first step for several graph mining algorithms, graph clustering is an active research area and numerous algorithms have been developed for this problem. Since our objective is to minimize *leakage* (that is, the interaction between clusters) while maintaining clusters with roughly similar size, we have chosen an algorithm that solves the *balanced partitioning problem*. While this problem is computationally hard even for small instances, we seek to solve this problem in a distributed manner for large-scale graphs. Next, we formally define the balanced partitioning problem, and then present our distributed algorithms for this problem. Consider a number k, and a graph $G(V, E)$ of n vertices with edge lengths (or edge weights), and node weights, and a real number $\alpha \geq 0$. We denote the edge lengths by $\ell : E \to R$, and the node weights by $w : V \to R$. Let $w(S)$ for any subset S of nodes be the total weight of nodes in S, i.e., $w(S) = \sum_{i \in S} w(i)$. A partition of nodes of G into k parts $\{V_i : i \in [k]\}$ is said to be α-balanced if and only if $(1 - \alpha)w(V)/k \leq w(V_i) \leq (1 + \alpha)w(V)/k$. In particular, a zero-balanced (or fully balanced) partition is one where all partitions have the same weight. The weight of the cut for this partitioning is the total sum of all edges whose endpoints fall in different parts:

$$\sum_{\substack{i \\ j<i}} \sum_{\substack{u \in V_i, v \in V_j: \\ (u,v) \in E(G)}} \ell(u, v).$$

Our goal in the balanced partitioning problem is to find an α-balanced partitioning whose cut size is (approximately) minimized. In most cases, we are not given a specific α as input, and we aim to find a partitioning that is as balanced as possible, i.e., with minimum α.

Input : Undirected graph $G(V, E)$ with node and edge weights and desired cluster size U
Output: A partition of G into clusters of size near U

Graph Contraction:
$S \leftarrow$ `IdentifySeedSetMR` (G, U);
$C \leftarrow$ `NaturalCutsMR` (G, S, U) ; `// Cut edges`
$CC \leftarrow$ `ConnectedComponentsMR` $(G(V, E - C))$;
$H \leftarrow$ `GraphContractionMR` (G, CC);
In Memory Merge:
$P \leftarrow$ `GraphAssembly` (H);
`OutputClusters` (P);

Algorithm 2: Phase 2 of the GeoCUTS algorithm.

Partitioning large scale geographic graphs. As our main contribution in this section, we present a distributed implementation of balanced partitioning by combining the idea of natural cuts [5, 6] and linearly embedding the nodes of a geo graph [2, 13, 16]. In particular we use Hilbert curve embedding of geo graphs due to its nice space-filling properties [13]. Following the structure of algorithms based on natural cuts [5, 6], our distributed algorithm iterates on a contraction stage and generates the output by applying a post-processing contraction stage, in which we *merge* parts of the contracted graph to compute the final partitioning.

Contraction stage. Inspired by prior work [6, 5], we propose and implement a distributed version of the contraction stage of balanced partitioning based on natural cuts. This stage consists of a sequence of four main steps: (i) identifying a seed set, (ii) finding a natural cut around each seed node, (iii) computing connected components of the graph after removing edges of the natural cuts, and (iv) finally contracting nodes in each connected component to one node and computing an updated smaller graph with new node weights and edge weights. The main feature of our distributed implementation is that all of the natural cuts are computed in parallel, and the graph contraction based on these cuts also happens in a distributed manner. More specifically, the contraction stage of the algorithm is as follows:

1. **Identify a seed set.** Here, the goal is to identify a set of k seed nodes S from which we compute natural cuts. We follow the following strategy for computing this set S: embed nodes of the graph into a line using the Hilbert curve, and divide the line into k pieces, each with almost the same number of nodes. Then, output a node close to the center of each piece of the Hilbert curve embedding. This method ensures that the set of seed nodes are spread uniformly across different parts, and thus natural cuts cover different parts of the graph.

2. **Parallel computation of natural cuts around seed nodes.** After selecting the seed set, we compute a "natural cut" around each seed node in parallel. Consider a seed node v, and let $Q = \frac{w(V)}{k}$, i.e., Q is the maximum weight of a cluster in an α-balanced partitioning. The idea is first to compute a core $C(v)$ around node v by performing a BFS around node v until we cover $Q/10$ nodes. We contract $C(v)$ to a node s. Then we continue the BFS until the total size of the neighborhood reaches Q and form a graph $G'(v)$ around node v. We take the rest of the graph $G \backslash G'(v)$ and contract it to one node, denoted by t. Finally, we compute a minimum (s, t)-cut in this graph $G'(v)$. We call this (s, t)-cut a *natural cut* around seed node v. A desirable property of this cut is that it has less than Q total weight on its nodes, and can be used as a building block for computing parts of a balanced partitioning. Note that we compute all these natural cuts in parallel in a distributed manner by applying a MapReduce framework, uploading the graph in a distributed hash-table service, and accessing the neighborhood of nodes via a read-only service [2, 12].

3. **Distributed connected components.** As the next part of the contraction stage, we remove all edges of the graph that appear in at least one of the natural cuts computed around the seed nodes, and

then compute connected components in the remaining graph (after removing those edges). We apply a distributed implementation of connected components that employs a distributed hash-table service, and that has been shown to be effective and scalable in practice [12, 14].

4. **Contracting nodes of each connected component.** After computing connected components, we can easily construct a contracted graph as follows: we put a node u_i for each connected component T_i with node weight $w(T_i)$, i.e., the sum of the weight of nodes in T_i. We also set the weight of the edge between two nodes u_i and u_j to $w(u_i, u_j)$.

Merging stage. If the size of the contracted graph is large, we iterate on the contraction stage until the size of the graph is small enough to fit in memory.[1] When the contracted graph fits in memory, we can produce an output by applying any in-memory heuristic for balanced partitioning of graphs with node weights. The algorithm that we employ at this stage is similar to the greedy assembly algorithm proposed in [5, 6]. We refer for details of this algorithm to [5, 6].

4 Statistical evaluation

In order to evaluate the quality of the clusters, we define a set of metrics to measure how well a given partitioning supports the purpose of experimentation. The goal in designing any experimental methodology is to maximize statistical power. While the optimal partitioning may depend on the treatment assigned, we require our clusters to support a class of experiments, the effects of which are unknown in advance. Our metrics are thus designed to be treatment-independent and based solely on the structure of the partitions.

4.1 Cluster quality metric for leakage

One of our chief considerations in partitioning is leakage. The reason is that the fidelity of a geo region as a unit of experimentation depends upon it. If we experiment on a leaky region, users therein will receive a diluted dose of treatment. Let a_{ij} represent the number of queries performed by User i in Region j. If we treat only Cluster k, the fraction of treatment received by User i is given by

$$\frac{a_{ik}}{\sum_j a_{ij}}.$$

We do not know precisely how much a user will respond to partial treatment; no doubt it depends on the type of treatment. A moderate assumption is that the user will respond in a manner prorated to the treatment she receives. That is, if Cluster k receives one unit of treatment (assume for now it is the only cluster being treated), then User i will respond with $\frac{a_{ik}}{\sum_j a_{ij}}$ units of response.

Of course, our goal is not to have to measure the response of individual users. Instead, we wish to treat and measure the response of whole clusters. The response measured in Cluster k is simply the weighted average of all user responses in that cluster, the weight being the fraction of mass or activity the user comprises in the cluster. This leads to an overall response for Cluster k being

$$\sum_i \left(\frac{a_{ik}}{\sum_\ell a_{\ell k}}\right)\left(\frac{a_{ik}}{\sum_j a_{ij}}\right) = \sum_i \frac{a_{ik}^2}{(a_{\cdot k})(a_{i\cdot})} = Q_k,$$

where $a_{i\cdot}$ is the total number of queries issued by User i across all clusters and $a_{\cdot k}$ is the number of queries in Cluster k issued by all users. This is our definition of Q_k, the quality or *Q-metric* of Cluster k. It is the prorated response we would measure in Cluster k were we to apply to it one unit of treatment (and to no other cluster).

In reality, an experiment is likely to treat multiple clusters. Let p represent the fraction of clusters treated by the experiment. A user may therefore receive treatment outside of Cluster k. Furthermore this reduction

[1] For our data sets, we never needed to iterate on this stage, since after the first stage, the graph fits in memory.

in treatment leakage out of Cluster k also contributes to treatment leakage into control clusters. The result is that the average measured difference between response of a treatment cluster T and a control cluster C is given by the convex combination

$$(1-p)Q_T + pQ_C.$$

Thus the average experimental effect we measure using clusters as our randomized experimental unit is proportional to the average "quality" $\bar{Q}$ of the clusters. In this sense, $\bar{Q}$ measures the average attenuation of any experimental effect we induce on a given cluster. And this is proportional to the statistical power of our experiments.

The case for why the average Q-metric is meaningful in running experiments assumes a simple prorated response by users. Such a linear dose-response curve is not always justified. For many treatments, we expect no response at all until the dose goes above a threshold, or perhaps a sigmoid dose-response curve. Without knowledge of the specific treatment we cannot improve our model of how user response is attenuated due to leakage. However, we may choose the conservative route of only using clusters with very high Q-metric as experimental units. This approach trades some statistical power for robustness, and it is this approach that we will take in analyzing many of our experiments.

4.2 Cluster balance metric

While it is certainly possible to experiment with units of different size, there are practical reasons to wish for greater balance. One is that when assigning each cluster to a randomized treatment group with probability p, we hope to have somewhere near fraction p of the population in treatment. This is unlikely to occur exactly if the clusters are of unequal size. Let X_i represent the binary variable for assignment of Cluster i to treatment where $\Pr(X_i = 1) = p$. Let S be the size of the chosen treatment group. The expectation of S is $\mathbb{E}\left(\sum_i w_i X_i\right) = p$. However if we assign cluster to treatment independently

$$\text{var } S = \sum_i w_i^2 \text{var } X_i = \sum_i w_i^2 p(1-p) = p(1-p)|w|^2$$

where $|w|^2$ is the ℓ^2 norm of the cluster weights. Thus the average degree of imbalance is proportional to $|w|^2$. We can effect greater balance if we choose exactly k clusters for treatment out of a total of n clusters where $p = k/n$. In this case,

$$\text{var } S = \frac{k}{n}\left(\frac{n-k}{n-1}\right)\left(|w|^2 - \frac{1}{n}\right) \approx p(1-p)\left(|w|^2 - \frac{1}{n}\right),$$

the approximation being for $k \ll n$. Again, the ℓ^2 norm of cluster weights relates to imbalance in selecting the treatment group. This imbalance is minimized when every cluster has the same size. We refer to the quantity $\left(|w|^2 - \frac{1}{n}\right)$ as the *B-metric* of the clustering. It is equal to 0 if the clustering is perfectly balanced and is otherwise positive, with greater size indicating greater imbalance.

4.3 Effective number of clusters

As mentioned in the introduction, the goal of our clustering is to maximize experimental power. However this is not possible to do without some knowledge of the experimental effect being administered. For instance, if the clusters grouped together users whose responses to the experiment were uncorrelated, the size and number of clusters would not affect estimation variance, and hence nor experimental power. To see this, suppose we measure $\hat{\theta}$ as the average effect across a set of clusters. If w_i is the fraction of all users in Cluster i then under uncorrelation var $\hat{\theta}_i \propto \frac{1}{w_i}$. Thus,

$$\hat{\theta} = \sum_i w_i \hat{\theta}_i$$

$$\text{var } \hat{\theta} = \sum_i w_i^2 \text{var } \hat{\theta}_i = \sum_i w_i^2 \frac{\sigma^2}{w_i} = \sigma^2,$$

where σ^2 is a constant. Of course we cannot assume that the users in any cluster will have uncorrelated responses so this is just one extreme. On the opposite extreme, if every user within a cluster behaves identically then var $\hat{\theta}_i$ is constant, say σ^2. In this case,

$$\text{var } \hat{\theta} = \sum_i w_i^2 \text{var } \hat{\theta}_i = \sum_i w_i^2 \sigma^2 = \sigma^2 |w|^2.$$

We require knowledge of how correlation within the cluster varies with cluster size to determine the effect of clustering on statistical power, and this is not available without knowledge of the experimental treatment. We could take $|w|^2$ as proportional to estimation variance in the worst case. This could be exceedingly conservative, but the argument emphasizes the utility of measuring $|w|^2$ through the B-metric.

5 Empirical Results

In this section we evaluate our algorithm and compare it against alternative algorithms and baselines.

5.1 The dataset

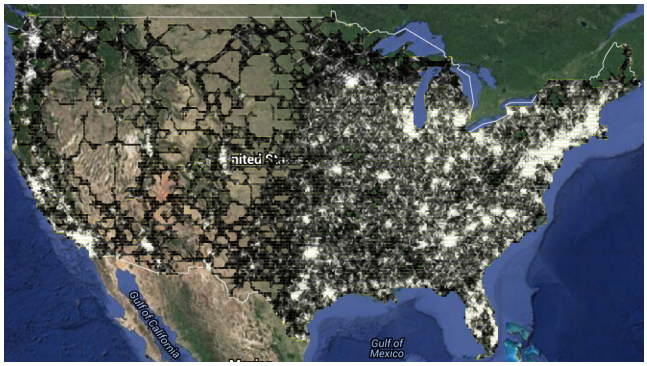

Figure 1: The graph built by the GeoCUTS algorithm, with nodes shown on corresponding locations of the US. White represents large edge weights (high traffic areas), while black represents low edge weights. Larger edge weights often do not match larger vertex weights, showing the difference between GeoCUTS and an algorithm simply measuring population density. Gaps in the colored regions represent locations for which no data is available. For example, in uninhabited regions such as deserts, data points come disproportionately from narrow strips corresponding to major roads.

A good partitioning algorithm minimizes the leakage introduced by the movement of users while keeping clusters roughly balanced. Hence, identifying movement trends is crucial in our algorithm. We wish to use clusters to run randomized experiments on user queries, hence a natural choice for constructing movement patterns is to use the approximate location of queries.

Our dataset consists of an anonymized set of randomly selected cookies. Each cookie is represented with a hashed value. The available data is an approximate location specified as a bounding box that covers the location at which a query was issued. The size of the bounding boxes is not uniform but all of them are large enough to contain locations of queries issued by a large number of distinct cookies. In some cases the boxes are so large that they cover an entire country or a continent. We remove such extreme cases from our dataset. Figure 3 shows the distribution of queries from the United States in terms of the size of their bounding boxes. This distribution varies markedly between countries and regions and tends to be skewed towards larger boxes in rural areas, where geo-location is less accurate and where the smaller number of queries also makes it necessary for bounding boxes to be larger in order to ensure anonymization. To build a graph, we form a grid on the geographical area we wish to partition. Each graph node corresponds to a grid

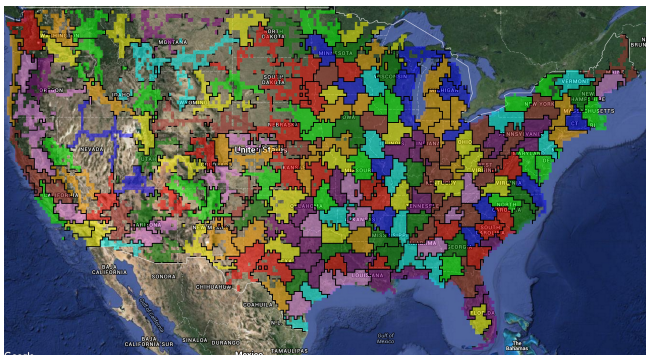

Figure 2: The GeoCUTS algorithm applied to user queries from the United States. The algorithm automatically identifies metropolitan areas, correctly predicting, for example, that the Bay Area includes San Francisco, Berkeley, and Palo Alto, but not Sacramento.

cell. We establish edges between two nodes based on the number of movements between their corresponding cells, that is, when the same cookie issues queries in both cells. Hence, for each query we only need to identify the cell it is issued from; for this purpose, we assume that the unknown query location is issued at the center of the known bounding box. To ensure that inaccuracies in estimating positions do not negatively impact our algorithm, the grid cells should be large relative to the typical sizes of the bounding boxes. We will take this assumption into account when we discuss the granularity of the grid cells. We build the graph

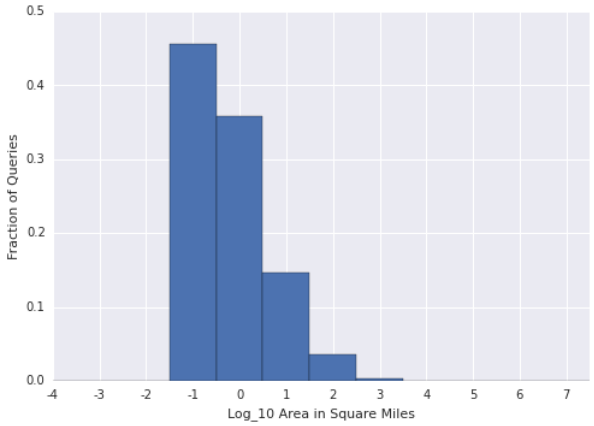

Figure 3: The distribution of sizes of bounding boxes for queries within the United States. Most bounding boxes are small, with an area of no more than 10 square miles, while about 4% have an area of 100 square miles or greater.

based on 28 days of cookie movements between grid cells with respect to user queries. The approximate locations of a very large number of queries are considered to model movement trends. Due to the scale of our datasets, while we have a very rough position for each individual cookie's queries, the aggregate movement patterns are quite accurate.

In the rest of this section, we evaluate the performance of the GeoCUTS algorithm on different datasets, for each of which a separate graph was built and clustered. Figures 2 and 4 show the clusters generated for data from the United States and France, respectively, and Figure 1 shows the GeoCUTS graph before clustering.

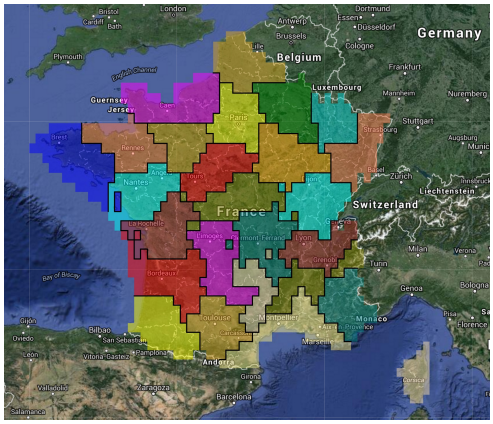

Figure 4: The GeoCUTS algorithm applied to user queries from France. It correctly identifies metropolitan areas such as Paris, Bordeaux, and Lyon, and regions such as Alsace and Normandy.

5.2 Mobility

Cookies that do not move are not interesting for our problem. In an extreme scenario where all cookies are stationary and issue queries from one location only, any arbitrary clustering algorithm performs perfectly well in terms of leakage. Hence for our first dataset we consider only *highly mobile* cookies: cookies that issue queries in at least two different cells of the grid.

We also note that some cookies are churned rather shortly after they are created. In practice, multiple cookies represent the same user over the period of our analysis. Hence, cookies with low query frequency under-represent the true movement. Query bounding boxes are in fact samples from the actual movement path, and a small number of samples is not enough to reconstruct the path. Hence for our second dataset we consider *highly active* cookies only: cookies with a query in more than 10 out of the 28 days.

Some highly active users may still have a limited movement and issue queries from the same geographical area. On the other hand, while some highly mobile users may issue fewer queries, they tend to move over a wider range. Therefore we expect leakage to be higher for highly mobile users, which is validated in our experiments.

For each of the United States and France, we collect two datasets - one for highly mobile and one for highly active cookies - and form a graph for each dataset. Both graphs have the same number of nodes (e.g. about 11,000 for the US for grid size 0.25 degrees). Unless otherwise specified, node and edge weights are log-normalized - in a subsequent section we will compare various normalization methods.

5.3 Comparison against other clusterings

We previously defined Q-metric to measure leakage. GeoCUTS partitions the map in a way that reduces cut size. Clusters with lower values of cut size tend to have higher values of Q-metric. In this section we compare GeoCUTS with other alternatives in terms of Q-metric and cut size.

The average and query-weighted average of the Q-metric for each clustering algorithm are listed in Table 1(a). According to every metric, GeoCUTS beats the baseline clustering corresponding to a coarse grid of regions. Where applicable, GeoCUTS and DMAs perform similarly well. It is important to note that in every evaluation, we compared only clusterings with similar numbers of clusters. Thus, in constructing the grid baseline, we picked the coarseness of the grid so that the number of regions in the grid approximated the number of clusters formed by GeoCUTS. We also used ~ 200 clusters for GeoCUTS in the US, in order to provide an effective comparison with DMAs.

The fraction of clusters and queries for different lower bounds of Q-metric are shown in Table 1(b). For example 80% of queries in the highly active set are issued from minimum-cut clusters with a Q-metric of at

(a)		**GeoCUTS**		**DMA**		**Grid**	
		Avg	Query-weighted avg	Avg	Query-weighted ag	Avg	Query-weighted avg
US	Highly Active	86.9%	92.1%	88.1%	91.7%	84.0%	90.5%
	Highly Mobile	79.0%	85.4%	80.1%	85.0%	76.0%	80.7%
France	Highly Active	84.2%	88.8%	-	-	82.8%	86.3%
	Highly Mobile	74.4%	78.6%	-	-	75.0%	77.2%

(b)		**GeoCUTS**			**DMA**			**Grid**		
		≥ 0.75	≥ 0.8	≥ 0.85	≥ 0.75	≥ 0.8	≥ 0.85	≥ 0.75	≥ 0.8	≥ 0.85
US	Highly Active	99.8%	99.0%	96.6%	100.0%	99.6%	97.9%	99.8%	98.8%	94.3%
	Highly Mobile	96.1%	86.4%	51.7%	95.3%	81.4%	48.5%	93.6%	60.0%	10.5%
France	Highly Active	99.8%	97.1%	79.7%	-	-	-	99.7%	88.8%	64.0%
	Highly Mobile	78.0%	42.3%	11.3%	-	-	-	74.4%	24.3%	5.0%

Table 1: (a) Average and query-weighted average of Q-metric, (b) Percentage of queries from clusters with a Q-metric of at least $x\%$. ~ 200 clusters were used for the US and ~ 50 for France.

		GeoCUTS			**Grid**		
		≥ 0.75	≥ 0.8	≥ 0.85	≥ 0.75	≥ 0.8	≥ 0.85
~ 25 clusters	Highly Active	100%	100%	94.3%	99.6%	99.1%	69.1%
	Highly Mobile	93.6%	55.5%	9.5%	83.0%	30.1%	0%
~ 50 clusters	Highly Active	99.8%	97.1%	79.7%	99.7%	88.8%	64.0%
	Highly Mobile	78.0%	42.3%	11.3%	74.4%	24.3%	5.0%

Table 2: Percentage of queries from clusters with a Q-metric $\geq x\%$ for different numbers of clusters in France.

least 0.8. For the grid, 62% of queries are issued from clusters with a Q-metric of at least 0.8. As already noted, highly mobile graphs are more challenging to partition compared to highly active graphs. The data indicates that the gap between our algorithm and baseline is larger for highly mobile graphs.

In Table 1, Q-metric values for DMAs are shown only for the US as they are undefined for other countries such as France. For both highly active and highly mobile graphs, GeoCUTS performs comparably to DMAs. Unlike DMAs, GeoCUTS offers flexibility in the number of clusters and functioning outside the US.

	GeoCUTS	DMA	Grid
US Highly Active	0.015	0.015	0.015
US Highly Mobile	0.018	0.017	0.013
France Highly Active	0.111	-	0.115

Table 3: B-metrics compared to other clusterings.

While the Q-metric quantifies the leakage, we still need to compare the clustering algorithms in terms of balance. An algorithm that produces highly unbalanced clusters may outperform other alternatives if only the Q-metric is considered. For example, if we partitioned the US into 200 clusters where 199 of them were in Alaska and one cluster covered the rest of the country, we would obtain an almost perfect Q-metric as relatively few users would cross between clusters. Obviously, however, such a clustering would not be useful for our applications. We compare B-metrics in Table 3. The results indicate that GeoCUTS performs equally well in terms of balance as the alternatives for highly active graphs, and performs slightly better for the highly mobile graph. In summary, while we perform better in terms of leakage, we do not compromise balance.

We next compare the cut-sizes generated by GeoCUTS to DMA and grid clusters as baselines. Table 4 shows the cut size comparison on the US log-normalized graph with grid size 0.25. It is clear that the GeoCUTS algorithm produces much better cut sizes compared to DMA and grid partitioning. Finally, we

	GeoCUTS	DMA	Grid	LE	Hilbert
Highly Active	**3.5%**	6.7%	14.8%	4.2%	7.3%
Highly Mobile	**3.8%**	7.2%	14.0%	4.4%	7.4%

Table 4: Cut size comparison against different clustering algorithms. "Grid" denotes the grid partition, "LE" denotes the Linear Embedding algorithm [2].

compare GeoCUTS results with Linear Embedding [2]. Since GeoCUTS itself uses Hilbert Curve for seed-selection we also compared it with partitions generated along a Hilbert Curve [13]. Table 4 compares the performance of alternative clustering algorithms; again GeoCUTS algorithm outperforms others on the same graph.

5.4 Normalization

We compared the performance of the GeoCUTS algorithm for different types of normalization during the graph-building phase. Specifically, we built graphs over US queries using logarithmic normalization of both vertices and edges, square root normalization, and also no normalization step at all. As expected, a stronger normalization was associated with better Q-metrics but worse B-metrics, demonstrating that normalization may be seen as mediating the tradeoff between diminished leakage and increased balance. The results of our comparison are shown in Table 5(a).

5.5 Cell width

Finally, we compare the performance of GeoCUTS across varying coarsenesses of location discretization. The majority of our experiments are conducted with locations rounded to the nearest 0.25 degree (latitude and longitude). Here we compare log-normalized graphs in which locations are discretized to 0.1, 0.25, and 0.5 degrees. It is worth noting that all of these sizes are considerably larger than the side length of a typical bounding box for location data. As expected, the coarsest discretization of 0.5 performs the worst in Q-metric and best in B-metric, as coarser discretization enforces balance but reduces the ability to decrease leakage. The results of our comparison are shown in Table 5(b).

(a)	Log	Square root	None
Q-metric, Highly Active	**0.9205**	0.8809	0.8395
Q-metric, Highly Mobile	**0.8539**	0.8069	0.7645
B-metric, Highly Active	0.0165	0.0047	**0.0006**
B-metric, Highly Mobile	0.0182	0.0053	**0.0011**

(b)	0.1	0.25	0.5
Q-metric, Highly Active	0.9157	**0.9205**	0.8580
Q-metric, Highly Mobile	0.8471	**0.8539**	0.7813
B-metric, Highly Active	0.0157	0.0165	**0.0115**
B-metric, Highly Mobile	0.0175	0.0182	**0.0132**

Table 5: Comparison of weighted average Q-metrics and B-metrics across (a) varying normalizations, (b) varying coarsenesses of location discretization.

6 Conclusion

We have presented an algorithm, GeoCUTS, designed for optimizing the assignment of regions to user queries. These regions can be used in running A/B bucket experiments to measure user response under treatment. Clustering users based on geographic region offers two major advantages: 1) assigning identical treatments

to different browser cookies of the same user, and 2) decreasing the probability of interference if users with different treatments interact. Unlike existing systems, GeoCUTS can be run in any region of the world and for any number of clusters. Furthermore, the algorithm is based on user traffic as measured by queries.

We have presented a system of metrics for statistically evaluating the performance of clustering algorithms for A/B bucket testing. According to these and prior metrics, GeoCUTS performs significantly above baseline, and comparably to the state-of-the-art hand-designed clustering given by DMAs. We have also analyzed the parameter settings used in the GeoCUTS algorithm, showing how normalization of node and edge weights allows us to interpolate between the extremes of highly balanced clusters and minimal leakage.

References

[1] K. Andreev and H. Räcke. Balanced graph partitioning. *Theory Comp. Sys.*, 39(6):929–939, 2006.

[2] K. Aydin, M. Bateni, and V. Mirrokni. Distributed balanced partitioning via linear embedding. In *WSDM 2016*.

[3] L. Backstrom and J. Kleinberg. Network bucket testing. In *WWW 2011*.

[4] L. Backstrom, E. Sun, and C. Marlow. Find me if you can: improving geographical prediction with social and spatial proximity. In *WWW 2010*.

[5] D. Delling, A. V. Goldberg, I. Razenshteyn, and R. F. F. Werneck. Exact combinatorial branch-and-bound for graph bisection. In *ALENEX 2012*.

[6] D. Delling, A. V. Goldberg, I. Razenshteyn, and R. F. F. Werneck. Graph partitioning with natural cuts. In *IPDPS 2011*.

[7] D. Eckles, B. Karrer, and J. Ugander. Design and analysis of experiments in networks: Reducing bias from interference. *Preprint arXiv:1404.7530*, 2014.

[8] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.

[9] H. Gui, Y. Xu, A. Bhasin, and J. Han. Network A/B testing: From sampling to estimation. In *WWW 2015*.

[10] H. Hohnhold, D. O'Brien, and D. Tang. Focusing on the long-term: It's good for users and business. In *KDD 2015*.

[11] L. Katzir, E. Liberty, and O. Somekh. Framework and algorithms for network bucket testing. In *WWW 2012*.

[12] R. Kiveris, S. Lattanzi, V. Mirrokni, V. Rastogi, and S. Vassilvitskii. Connected components in mapreduce and beyond. In *SOCC 2014*.

[13] B. Moon, H. V. Jagadish, C. Faloutsos, and J. H. Saltz. Analysis of the clustering properties of the Hilbert space-filling curve. *IEEE Trans. on Knowledge & Data Eng.*, 13, 2001.

[14] V. Rastogi, A. Machanavajjhala, L. Chitnis, and A. D. Sarma. Finding connected components in map-reduce in logarithmic rounds. In *ICDE 2013*.

[15] I. Stanton and G. Kliot. Streaming graph partitioning for large distributed graphs. In *KDD 2012*.

[16] R. G. Strongin and Y. D. Sergeyev. *Global optimization with non-convex constraints: sequential and parallel algorithms*. Nonconvex Optimization and Its Applications. Kluwer academic publishers, 2000.

[17] D. Tang, A. Agarwal, D. O'Brien, and M. Meyer. Overlapping experiment infrastructure: More, better, faster experimentation. In *KDD 2010*.

[18] C. Tsourakakis, C. Gkantsidis, B. Radunovic, and M. Vojnovic. Fennel: Streaming graph partitioning for massive scale graphs. In *WSDM 2014*.

[19] J. Ugander and L. Backstrom. Balanced label propagation for partitioning massive graphs. In *WSDM 2013*.

[20] J. Ugander, B. Karrer, L. Backstrom, and J. Kleinberg. Graph cluster randomization: Network exposure to multiple universes. In *KDD 2013*.

[21] J. Vaver and J. Koehler. Periodic measurement of advertising effectiveness using multiple-test-period geo experiments. Technical report, Google Inc., 2012.

[22] X. Zhu and Z. Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical report, Citeseer, 2002.